manuscript No.
(will be inserted by the editor)

On general sampling schemes for Particle Markov chain Monte Carlo methods

Eduardo F. Mendes · **Christopher K. Carter** · **Robert Kohn**

arXiv:1401.1667v2 [stat.CO] 12 Jan 2015

Received: date / Accepted: date

Abstract Particle Markov Chain Monte Carlo methods (PMCMC) (Andrieu et al. 2010) are used to carry out inference in non-linear and non-Gaussian state space models, where the posterior density of the states is approximated using particles. Current approaches usually carry out Bayesian inference using a particle Marginal Metropolis-Hastings algorithm, a particle Gibbs sampler, or a particle Metropolis within Gibbs sampler. Our article gives a general approach for constructing sampling schemes that converge to the target distributions given in Andrieu et al. (2010) and Olsson and Ryden (2011). Our approach shows how the three ways of generating variables mentioned above can be combined flexibly. The advantage of our general approach is that the sampling scheme can be tailored to obtain good results for different applications. We investigate the properties of the general sampling scheme, including conditions for uniform convergence to the posterior. We illustrate our methods with examples of state space models where one group of parameters can be generated in a straightforward manner in a particle Gibbs step by conditioning on the states, but where it is cumbersome and inefficient to generate such parameters when the states are integrated out. Conversely, it may be necessary to generate a second group of parameters without conditioning on the states because of the high dependence between such parameters and the states. A particularly important case that is illustrated by our examples occurs when there are many unknown parameters. In this case it is usually highly inefficient to generate all the parameters without conditioning on the states and it is preferable to generate as many parameters as possible using particle Gibbs and only generate those parameters that are highly correlated with the states using particle Marginal Metropolis-Hastings.

Keywords Gibbs sampler · Metropolis-Hastings method · non-Gaussian · Particle filter · State space models

1 Introduction

Our article deals with statistical inference for both the unobserved states and the parameters in a class of state space models. Its main goal is to give a general approach to constructing

Eduardo F. Mendes, School of Economics, University of New South Wales · Christopher K. Carter, School of Economics, University of New South Wales, Tel.: +612-9385-9696, Fax: +612-9313-6337, E-mail: chris.carter@unsw.edu.au · Robert Kohn, School of Economics, University of New South Wales

sampling schemes that converge to the posterior distribution of the states and the parameters. The sampling schemes generate particles as auxiliary variables. This work extends the methods proposed by Andrieu et al. (2010), Lindsten and Schön (2012), Lindsten et al. (2014) and Olsson and Ryden (2011).

If the state space model is Gaussian, or Gaussian conditional on a set of latent state variables, then there are well established MCMC methods for carrying out Bayesian inference e.g. Carter and Kohn (1994), Frühwirth-Schnatter (1994), Gerlach et al. (2000) and Frühwirth-Schnatter (2006). However, in many cases it is not possible to use such conditionally Gaussian formulations, and it is necessary to use other methods, e.g. single site sampling of the states as proposed by Carlin et al. (1992), or block sampling of the states conditional on the parameters as proposed by Shephard and Pitt (1997). Such methods can be very inefficient, and often cannot be applied in a general way as they require a number of tuning parameters to be specified, such as the block size.

The bootstrap (standard) particle filter was introduced by Gordon et al. (1993), and independently by Kitagawa (1996), for online filtering of general state space models, with the parameters assumed known. The bootstrap particle filter just requires that it is possible to generate from the state transition equation and evaluate the observation equation. Pitt and Shephard (1999) introduce the class of auxiliary particle filters which can be much more efficient than the bootstrap filter, especially when the signal to noise ratio is high. There have been a number of approaches that use the particle filter for Bayesian inference on the parameters. Liu and West (2001) perform online inference including the unknown parameter vector as part of the states. Storvik (2002) and Polson et al. (2008) use a sequential updating scheme that relies on conjugacy to update the parameters.

Andrieu et al. (2010) introduce two particle MCMC methods. The first is particle marginal Metropolis Hastings (PMMH), where the parameters are generated with the states integrated out. The second is particle Gibbs (PG), which generates the parameters given the states. They show that the augmented density targeted by this algorithm has the joint posterior density of the parameters and states as a marginal density. Andrieu et al. (2010) and Andrieu and Roberts (2009) show that the law of the marginal sequence of parameters and states, sampled using either PG or PMMH, converges to the true posterior as the number of iterations increase. Both particle MCMC methods have been focus of recent research. Olsson and Ryden (2011) and Lindsten and Schön (2012) use *backward simulation* (Godsill et al. 2004) for sampling the state vector, instead of Kitagawa's *ancestral tracing*. These authors derive the augmented target distribution and find that, in simulation studies, their method performs better than the PMMH or its PG counterpart using *ancestral tracing*. Chopin and Singh (2013) show that *backward simulation* is more efficient than *ancestral tracing* in a PG setting, as well as being robust to the resampling method used (e.g., multinomial, residual, stratified). Lindsten and Schön (2012) extend the PG sampler to a particle Metropolis within Gibbs (PMwG) sampler to deal with the case where the parameters cannot be generated exactly conditional on the states.

In a separate line of research Pitt et al. (2012), Doucet et al. (2015) and Sherlock et al. (2015), show that the best tradeoff between computational cost and integrated autocorrelation time is attained when the standard deviation of the estimated log-likelihood is between about 1.0 and 1.8, the exact value depending on the quality of the proposal in the ideal or infinite particle case. The better the proposal, the smaller should the standard deviation be. For PG, it is unclear what role the standard deviation of the estimated likelihood plays in the choice of the number of particles.

MCMC using particle methods has two attractions. First, essentially the same code can be used for a variety of models. All that is usually necessary is to code up a new ver-

sion of the observation equation and possibly the state transition equation, as well as code for generating the states. Furthermore, any Metropolis within Gibbs step used in standard MCMC simulation can be used in the same way for particle based MCMC. Second, it is relatively straightforward to integrate out the states. However, since particle methods are computationally intensive, it is important to implement them efficiently. Therefore, using only Metropolis-Hastings or Gibbs samplers may be undesirable. We know from the literature on Gaussian and conditionally Gaussian state space models that confining MCMC for state space models to Gibbs sampling or Metropolis-Hastings sampling can result in inefficient or even degenerate sampling. See, for example, Kim et al. (1998) who show for a stochastic volatility model that generating the states conditional on the parameters and the parameters conditional on the states can result in a highly inefficient sampler. See also Carter and Kohn (1996) and Gerlach et al. (2000) who demonstrate using a signal plus noise model that a Gibbs sampler for the states and indicator variables for the structural breaks produces a degenerate sampler. A natural solution is to combine Gibbs and Metropolis-Hastings samplers. A second reason for combining Metropolis-Hastings samplers with Gibbs samplers for particle MCMC is that we may wish to consider generating as many of the parameters by PG as practical. This is because it is usually difficult to obtain good proposals for the parameters in the particle Metropolis-Hastings step because the likelihood is estimated and it is usually not possible to obtain the derivatives of the likelihood. This is unlike standard MCMC where the likelihood and its derivatives can be computed, at least numerically. Section 4 gives a fuller discussion of this issue.

Our work extends the particle MCMC framework introduced in Andrieu et al. (2010), Lindsten and Schön (2012) and Lindsten et al. (2014) to situations where using just PMMH or just PG is impossible or inefficient. We derive a general particle sampler on the same augmented space as the PG and PMMH samplers, in which some parameters are sampled conditionally on the states and the remaining parameters are sampled with the states integrated out. We show that the general sampler targets the same augmented density as the PG or PMMH samplers and that the Markov chain generated by the algorithm is uniformly ergodic, given regularity conditions. It implies that the marginal law of the Markov chain generated by n^{th} iteration of the algorithm, converges to the posterior density function geometrically fast, uniformly on the starting value, as $n \to \infty$.

The method is introduced in a simple Markov state-space setting, with a possibly adapted proposal density and using *ancestral tracing*. The proposed method can be modified in a straightforward way to a non-Markov case (the state at time t depends on the states at times $1,\ldots,t-1$), auxiliary particle filters and *backward simulation*. Although these modifications are useful in applications, we choose to sacrifice generality in favor of a more accessible presentation. The proofs may also be modified using arguments found in Olsson and Ryden (2011), and the same results hold. We conjecture it is possible to show geometric ergodicity, instead of uniform ergodicity, using arguments similar to those in Propositions 7 and 9 of Andrieu and Vihola (2012), but both the notation and algebra are more involved and it is not the main purpose of this paper.

The paper is organized as follows. Section 2 introduces the basic concepts and notation used throughout the paper. Section 3 discusses the general particle samplers. Section 4 reports the simulated and real examples and Section 5 contains some discussion. Appendix A shows that the general sampler is ergodic, under regularity conditions. Appendix B discusses the extension to *backward simulation* and Appendix C describes the cubic spline smoothness prior used in the second example.

2 Preliminaries

This section introduces the models and the main algorithms to generate Metropolis-Hastings proposals and to generate from useful conditional distributions. These are the main components that are used in the sampling schemes in Section 3. There are strong relationships between these components because the target distribution in Section 2.3 is on an augmented space that includes the particles from the sequential Monte Carlo algorithm described in Section 2.2. It is chosen to have marginal distributions that match the state space model in Section 2.1 and to include terms from the sequential Monte Carlo algorithm. The algorithm in Section 2.4 generates from an exact conditional distribution and is used for the Gibbs sampling steps described in Section 3.

2.1 State Space Model

Define $\mathbb{N}$ as the set of positive integers and let $\{X_t\}_{t\in\mathbb{N}}$ and $\{Y_t\}_{t\in\mathbb{N}}$ denote $\mathscr{X}$-valued and $\mathscr{Y}$-valued stochastic processes, where $\{X_t\}_{t\in\mathbb{N}}$ is a latent Markov process with initial density $\mu_\theta(x)$ and transition density $f_\theta(x'|x)$, i.e.,

$$X_1 \sim \mu_\theta(\cdot) \quad \text{and} \quad X_{t+1}|(X_t = x) \sim f_\theta(\cdot|x) \quad (t = 1,2,\dots).$$

The latent process $\{X_t\}_{t\in\mathbb{N}}$ is observed only through $\{Y_t\}_{t\in\mathbb{N}}$, whose value at time t depends on the value of the hidden state at time t, and is distributed according to $g_\theta(y|x)$:

$$Y_t|(X_t = x) \sim g_\theta(\cdot|x) \quad (t = 1,2,\dots).$$

The densities μ_θ, f_θ and g_θ are indexed by a parameter vector $\theta \in \Theta$, where Θ is an open subset of $\mathbb{R}^{d_\theta}$, and all densities are with respect to suitable dominating measures, denoted dx and dy. The dominating measure is frequently taken to be the Lebesgue measure if $\mathscr{X} \in \mathscr{B}(\mathbb{R}^{d_x})$ and $\mathscr{Y} \in \mathscr{B}(\mathbb{R}^{d_y})$, where $\mathscr{B}(A)$ is the Borel σ-algebra generated by the set A.

We use the colon notation for collections of random variables, i.e., for integers $r \le s$, $a_{r:s} = (a_r,\dots,a_s)$, $a_t^{r:s} = (a_t^r,\dots,a_t^s)$ and for $t \le u$, $a_{r:s}^{t:u} = (a_r^{t:u},\dots,a_s^{t:u})$. The joint probability density function of $(x_{1:T}, y_{1:T})\}$ is

$$p(x_{1:T}, y_{1:T}|\theta) = \mu_\theta(x_1) g_\theta(y_1|x_1) \prod_{t=2}^{T} f_\theta(x_t|x_{t-1})\, g_\theta(y_t|x_t).$$

We denote the likelihood by $Z_1(\theta) = p(y_1|\theta)$ and $Z_t(\theta) = p(y_t|y_{1:t-1},\theta)$ for $t \ge 2$. The joint filtering density of $X_{1:t}$ is

$$p(x_{1:t}|y_{1:t},\theta) = \frac{p(x_{1:t}, y_{1:t}|\theta)}{Z_{1:t}(\theta)},$$

where

$$Z_{1:t}(\theta) = \prod_{i=1}^{t} Z_t(\theta).$$

The filtering density of θ and $X_{1:T}$ can also be factorized as

$$p(x_{1:T},\theta|y_{1:T}) = \frac{p(x_{1:T}, y_{1:T}|\theta)p(\theta)}{\overline{Z}_{1:T}},$$

where $\overline{Z}_{1:T} = \int_\Theta Z_{1:T}(\theta)\, p(\theta)\, \mathrm{d}\theta = p(y_{1:T})$. This factorization is used in the particle Monte Carlo Markov chain algorithms.

2.2 Sequential Monte Carlo

This involves constructing a particulate approximation using a sequential Monte Carlo (or a particle filter) method. The idea is to approximate the joint filtering densities $\{p(x_t|y_{1:t},\theta) : t = 1,2,\dots\}$ sequentially, using particles, i.e., weighted samples, $(x_t^{1:N}, \overrightarrow{w}_t^{1:N})$, drawn from auxiliary distributions m_t^θ. It requires the specification of *importance densities* $m_1^\theta(x_1) := m_1(x_1|Y_1 = y_1,\theta)$ and $m_t^\theta(x_t|x_{t-1}) := m_t(x_t|X_{t-1} = x_{t-1}, Y_{1:t} = y_{1:t},\theta)$, and a resampling scheme $\mathscr{M}(a_{t-1}^{1:N}|\overrightarrow{w}_{t-1}^{1:N})$, where each $a_{t-1}^i = k$ indexes a particle in $(x_{t-1}^{1:N}, \overrightarrow{w}_{t-1}^{1:N})$, and is sampled with probability $\overrightarrow{w}_{t-1}^k$. We refer to Doucet et al. (2000), Van Der Merwe et al. (2001), and Guo et al. (2005) for the choice of importance densities and Douc and Cappé (2005) for a comparison between resampling schemes.

The Sequential Monte Carlo algorithm used here is the same one as in Andrieu et al. (2010) and is defined as follows.

Algorithm 1 (Sequential Monte Carlo)

1. *For $t = 1$:*
   (a) *Sample X_1^i from $m_1^\theta(x)\mathrm{d}x$, for $i = 1,\dots,N$*
   (b) *Calculate the importance weights*

$$w_1^i = \frac{\mu_\theta(x_1^i)\, g_\theta(y_1|x_1^i)}{m_1^\theta(x_1^i)} \quad (i = 1,\dots,N),$$

   and normalize them to obtain $\overrightarrow{w}_1^{1:N}$.

2. *For $t = 2,3,\dots$:*
   (a) *Sample the ancestral indices $A_{t-1}^{1:N} \sim \mathscr{M}\left(a^{1:N}|\overrightarrow{w}_{t-1}^{1:N}\right)$*
   (b) *Sample X_t^i from $m_t^\theta\left(x|x_{t-1}^{a_{t-1}^i}\right)$, $i = 1,\dots,N$*
   (c) *Calculate the importance weights*

$$w_t^i = \frac{f_\theta\left(x_t^i|x_{t-1}^{a_{t-1}^i}\right) g_\theta\left(y_t|x_t^i\right)}{m_t^\theta\left(x_t^i|x_{t-1}^{a_{t-1}^i}\right)} \quad (i = 1,\dots,N)$$

   and normalize them to obtain $\overrightarrow{w}_t^{1:N}$.

We denote the vector of particles by

$$U_{1:T} := \left(X_1^{1:N},\dots,X_T^{1:N},A_1^{1:N},\dots,A_{T-1}^{1:N}\right)$$

and its sample space by $\mathscr{U} := \mathscr{X}^{TN} \times \mathbb{N}^{(T-1)N}$. The algorithm automatically provides an estimate of the likelihood $Z(u_{1:T},\theta) = \prod_{t=1}^T \left(N^{-1}\sum_{i=1}^N w_t^i\right)$. The joint distribution of the particles given the parameters is

$$\psi\left(u_{1:T}^{1:N}|\theta\right) := \prod_{i=1}^N m_1^\theta\left(x_1^i\right) \prod_{t=2}^T \left\{ \mathscr{M}(a_{t-1}^{1:N}|\overrightarrow{w}_{t-1}^{1:N}) \prod_{i=1}^N m_t^\theta\left(x_t^i|x_{t-1}^{a_{t-1}^i}\right)\right\}. \quad (1)$$

2.3 Target distribution

The key idea of particle MCMC methods is to construct a target distribution on an augmented space that includes the particles $U_{1:T}$ and has marginal distributions equal to $p(x_{1:T},\theta|y_{1:T})$. This section describes the target distribution from Andrieu et al. (2010). Later sections will describe current MCMC methods to sample from this distribution and hence sample from $p(x_{1:T},\theta|y_{1:T})$. Appendix B describes other choices of target distribution and how our results can be modified in a straightforward way to apply to these distributions.

The simplest way of sampling from the particulate approximation of $p(x_{1:T}|y_{1:T},\theta)$ is called *ancestral tracing*. It was introduced in Kitagawa (1996) and used in Andrieu et al. (2010) and consists of sampling one particle from the final particle filter. The method is equivalent to sampling an index $J=j$ with probability $\bar{w}_T^j$, tracing back its ancestral lineage $b_{1:T}^j$ ($b_T^j=j$ and $b_{t-1}^j=a_t^{b_t^j}$) and choosing the particle $x_{1:T}^j=(x_1^{b_1^j},\ldots,x_T^{b_T^j})$.

With some abuse of notation, for a vector a_t, denote $a_t^{(-k)}=\{(a_t^1,\ldots,a_t^{k-1},a_t^{k+1},\ldots,a_t^N)$, with obvious changes for $k\in\{1,N\}$, and denote

$$u_{1:T}^{(-j)}=\left\{x_1^{(-b_1^j)},\ldots,x_{T-1}^{(-b_{T-1}^j)},x_T^{(-j)},a_1^{(-b_1^1)},\ldots,a_{T-1}^{(-b_{T-1}^j)}\right\}.$$

It will simplify the notation to sometimes use the following one-to-one transformation

$$(u_{1:T},j)\leftrightarrow\left\{x_{1:T}^j,b_{1:T-1}^j,j,u_{1:T}^{(-j)}\right\}$$

We will switch between the two transformations and use whichever is more convenient and note that the right hand expression will sometimes be written as $\left\{x_{1:T},b_{1:T-1},j,u_{1:T}^{(-j)}\right\}$ without ambiguity. We make the following assumptions from (Andrieu et al. 2010, Section 4.1).

Assumption 1 *The resampling scheme $\mathscr{M}(a_{t-1}^{1:N}|\bar{w}_{t-1}^{1:N})$ satisfies*

$$P(A_{t-1}^i=k|\bar{W}_{t-1}^{1:N}=\bar{w}_{t-1}^{1:N})=\bar{w}_{t-1}^k$$

and

$$\mathbb{E}[O_{t-1}^k|\bar{w}_{t-1}^{1:N}]=N\,\bar{w}_{t-1}^k,$$

for $t=2,\ldots,T$, where O_{t-1}^k is the number of offspring of particle k propagated at time t.

Assumption 2 *For any $\theta\in\Theta$, define*

$$\begin{aligned}\mathscr{S}_n(\theta)&=\{x_{1:n}\in\mathscr{X}^n:\,p(x_{1:n}|y_{1:n},\theta)>0\}\quad\text{for } n=1,\ldots,T,\\ \mathscr{Q}_n(\theta)&=\{x_{1:n}\in\mathscr{X}^n:\,p(x_{1:n-1}|y_{1:n-1},\theta)\,m_n^\theta(x_n|x_{n-1})>0\}\\ &\quad\text{for } n=1,\ldots,T,\end{aligned}$$

with the convention that $p(x_{1:0}|y_{1:0}):=1$ and $m_0^\theta(x_0|x_{1:0}):=1$, and $\mathscr{S}=\{\theta\in\Theta;\,p(\theta|y_{1:T})>0\}$. For any $\theta\in\mathscr{S}$, $\mathscr{S}_n(\theta)\subseteq\mathscr{Q}_n(\theta)$ for $n=1,\ldots,T$.

The target distribution from Andrieu et al. (2010) is

$$\begin{aligned}&\tilde{\pi}^N\left(x_{1:T},b_{1:T-1},j,u_{1:T}^{(-j)},\theta\right):=\\ &\frac{p(x_{1:T},\theta|y_{1:T})}{N^T}\frac{\psi(u_{1:T}|\theta)}{m_1^\theta\left(x_1^{b_1}\right)\,\prod_{t=2}^T\bar{w}_{t-1}^{a_{t-1}^i}m_t^\theta\left(x_t^{b_t}|x_{t-1}^{a_{t-1}^{b_t}}\right)}.\end{aligned}\tag{2}$$

From Andrieu et al. (2010), under Assumption 1, equation (2) has the following marginal distribution

$$\tilde{\pi}^N(x_{1:T}, b_{1:T-1}, j, \theta) = \frac{p(x_{1:T}, \theta | y_{1:T})}{N^T}$$

and hence

$$\tilde{\pi}^N(x_{1:T}, \theta) = p(x_{1:T}, \theta | y_{1:T}).$$

Assumption 2 ensures that $\tilde{\pi}^N(u_{1:T}|\theta)$ is absolutely continuous with respect to $\psi(u_{1:T}|\theta)$, as summarized in the following lemma.

Lemma 1

$$\tilde{\pi}^N(\cdot|\theta) \ll \psi(\cdot|\theta) \qquad \textit{for all } \theta \in \mathscr{S}.$$

The proof is straightforward and is omitted.

2.4 Conditional Sequential Monte Carlo

The particle Gibbs algorithm in Andrieu et al. (2010) uses exact conditional distributions to give a Gibbs sampler. If we use the *ancestral tracing* augmented distribution given in (2) then this includes the conditional distribution given by

$$\tilde{\pi}^N\left(u_{1:T}^{(-j)} | x_{1:T}^j, b_{1:T-1}^j, j, \theta\right)$$

which involves constructing the particulate approximation conditional on a pre-specified path. The *conditional sequential Monte Carlo* algorithm, introduced in Andrieu et al. (2010), is a sequential Monte Carlo algorithm in which a particle $X_{1:T}^J = (X_1^{B^J}, \ldots, X_T^{B_T^J})$, and the associated sequence of ancestral indices $B_{1:T-1}^J$ are kept unchanged. In other words, the conditional sequential Monte Carlo algorithm is a procedure that resamples all the particles and indices except for $U_{1:T}^J = (X_{1:T}^J, A_{1:T-1}^J) = (X_1^{B_1^J}, \ldots, X_T^{B_T^J}, B_1^J, \ldots, B_{T-1}^J)$. The conditional sequential Monte Carlo algorithm (as in Andrieu et al. (2010)), consistent with $(x_{1:T}^j, a_{1:T-1}^j, j)$, is described as follows.

Algorithm 2 (Conditional Sequential Monte Carlo)

1. fix $X_{1:T}^j = x_{1:T}^j$ and $A_{1:T-1}^j = b_{1:T-1}^j$.
2. For $t = 1$:
   (a) Sample X_1^i from $m_1^\theta(x)\mathrm{d}x$, for $i \in \{1, \ldots, N\} \setminus \{b_1^j\}$.
   (b) Calculate the importance weights

$$w_1^i = \frac{\mu_\theta(x_1^i)\, g_\theta(y_1|x_1^i)}{m_1^\theta(x_1^i)} \quad (i = 1, \ldots, N),$$

   and normalize them to obtain $\bar{w}_1^{1:N}$.
3. For $t = 2, \ldots, T$:
   (a) Sample the ancestral indices

$$A_{t-1}^{-(b_t^j)} \sim \mathscr{M}\left(a^{(-b_t^j)} | \bar{w}_{t-1}^{1:N}\right).$$

(b) Sample X_t^i from $m_t^\theta\left(x|x_{t-1}^{a_{t-1}^i}\right)\mathrm{d}x$, $i \in \{1,\ldots,N\} \setminus \{b_t^j\}$.

(c) Calculate the importance weights

$$w_t^i = \frac{f_\theta\left(x_t^i|x_{t-1}^{a_{t-1}^i}\right) g_\theta\left(y_t|x_t^i\right)}{m_t^\theta\left(x_t^i|x_{t-1}^{a_{t-1}^i}\right)} \quad (i=1,\ldots,N)$$

and normalized them to obtain $\overline{w}_t^{1:N}$.

Under assumption 1, it is straightforward to show that the algorithm samples from the target density of the random variable $U_{1:T}^{(-J)} = \left(X_1^{(-B_1^J)},\ldots,X_T^{(-B_T^J)},A_1^{(-B_2^J)},\ldots,A_{T-1}^{(-B_T^J)}\right)$, conditional on $U_{1:T}^J$ and index J given by

$$\tilde{\pi}^N\left(u_{1:T}^{(-j)}|x_{1:T},b_{1:T-1},j,\theta\right) = \frac{\psi\left(u_{1:T}|\theta\right)}{m_1^\theta\left(x_1^{b_1}\right)\ \prod_{t=2}^T \overline{w}_{t-1}^{a_{t-1}^i} m_t^\theta\left(x_t^{b_t}|x_{t-1}^{a_{t-1}^{b_t}}\right)};$$

see Andrieu et al. (2010) for details.

3 Sampling schemes

This section describes general sampling schemes using the distributions and algorithms from Section 2. We then discuss special cases including the PMMH algorithms and the PG algorithms of Andrieu et al. (2010).

3.1 General sampling scheme

This sampling scheme is suitable for state space models where some of the parameters can be generated exactly conditional on the state vectors, but other parameters must be generated using Metropolis-Hasting proposals. Let

$$\theta = (\theta_1,\ldots,\theta_p)$$

be a partition of the parameter vector into p components where each component may be a vector and let $0 \leq p_1 \leq p$. Let $\Theta = \Theta_1 \times \ldots \times \Theta_p$ be the corresponding partition of the parameter space. We will use the notation

$$\theta_{-i} = (\theta_1,\ldots,\theta_{i-1},\theta_{i+1},\ldots,\theta_p).$$

The following sampling scheme generates the parameters $\theta_1,\ldots,\theta_{p_1}$ using particle Marginal Metropolis-Hasting steps and the parameters $\theta_{p_1+1},\ldots,\theta_p$ using particle Gibbs or particle Metropolis within Gibbs steps. We call this a *general sampler*. To simplify the discussion we assume that both particle Marginal Metropolis-Hastings steps and particle Gibbs steps are used, i.e., $0 < p_1 < p$.

Sampling Scheme 1 (General Sampler) *Given initial values for $U_{1:T}$, J and θ, one iteration of the MCMC involves the following steps*
1. (PMMH sampling) For $i = 1, \ldots, p_1$
Step i:

(a) Sample $\theta_i^* \sim q_{i,1}(\cdot|U_{1:T}, J, \theta_{-i}, \theta_i)$.
(b) Sample $U_{1:T}^* \sim q_{i,2}(\cdot|U_{1:T}, J, \theta_{-i}, \theta_i^*)$.
(c) Sample J^* from $\tilde{\pi}^N(\cdot|U_{1:T}^*, \theta_{-i}, \theta_i^*)$.
(d) Accept the proposed values θ_i^*, $U_{1:T}^*$ and J^* with probability

$$\alpha_i(U_{1:T}, J, \theta_i; U_{1:T}^*, J^*, \theta_i^*|\theta_{-i}) = 1 \wedge \frac{\tilde{\pi}^N(U_{1:T}^*, \theta_i^*|\theta_{-i})}{\tilde{\pi}^N(U_{1:T}, \theta_i|\theta_{-i})} \frac{q_i(U_{1:T}, \theta_i|U_{1:T}^*, J^*, \theta_{-i}, \theta_i^*)}{q_i(U_{1:T}^*, \theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i)} \tag{3}$$

where

$$q_i(U_{1:T}^*, \theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i) = q_{i,1}(\theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i) q_{i,2}(U_{1:T}^*|U_{1:T}, J, \theta_{-i}, \theta_i^*).$$

2. (PG or PMwG sampling) For $i = p_1 + 1, \ldots, p$
Step i:

(a) Sample $\theta_i^* \sim q_i(\cdot|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i)$.
(b) Accept the proposed value θ_i^* with probability

$$\alpha_i\left\{\theta_i; \theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right\} = 1 \wedge \frac{\tilde{\pi}^N\left(\theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right)}{\tilde{\pi}^N\left(\theta_i|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right)} \times \frac{q_i(\theta_i|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i^*)}{q_i(\theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i)}. \tag{4}$$

3. Sample $U_{1:T}^{(-J)} \sim \tilde{\pi}^N(\cdot|X_{1:T}^J, B_{1:T-1}^J, J, \theta)$ using Algorithm 2.
4. Sample $J \sim \tilde{\pi}^N(\cdot|U_{1:T}, \theta)$.

Note that Parts 2, 3 and 4 are the same as the particle Gibbs sampler described in Andrieu et al. (2010) or the particle Metropolis within Gibbs sampler described in Lindsten and Schön (2012). Part 1 differs from the particle Marginal Metropolis-Hastings approach discussed in Andrieu et al. (2010) by a more general Metropolis-Hastings proposal. Part 1 also generates the variable J which selects the trajectory. This is necessary since J is used in Part 2.

A major computational cost of the algorithm is generating the particles p_1 times in Part 1 and running the Conditional SMC algorithm in Part 3. Hence there is a computational cost in using the general sampler compared to a particle Gibbs sampler. Similar comments apply to a blocked PMMH sampler.

Remark 1 *If the general Metropolis-Hastings proposals described in Tierney (1998) are used in Part 1, then it is possible to obtain the particle Gibbs or particle Metropolis within Gibbs steps in Part 2 as special cases. In this paper, we separate the cases as shown above so we can discuss the computational efficiency and the mixing rates of different sampling schemes. We have chosen the most commonly used methods to simplify the discussion and more general sampling schemes may be used in some applications.*

Remark 2 *Part 1 of the sampling scheme is a good choice for parameters θ_i where a good approximation for the distribution*

$$\tilde{\pi}^N(\theta_i|\theta_{-i}) = p(\theta_i|y_{1:T},\theta_{-i})$$

is available as a Metropolis-Hastings proposal, whereas Part 2 of the sampling scheme is more suited for parameters θ_i where it is possible to sample exactly from the distribution

$$\tilde{\pi}^N(\theta_i|X_{1:T},\theta_{-i}) = p(\theta_i|X_{1:T},y_{1:T},\theta_{-i})$$

or a good approximation is available as a Metropolis-Hastings proposal. See Lindsten and Schön (2012) for more discussion about the particle Metropolis-Hasting proposals in Part 2.

Remark 3 *It is possible to simplify Part 4 of the sampling scheme by not generating J. See van Dyk and Park (2008) for a discussion on methods to obtain partially collapsed Gibb samplers. We believe in this case, however, that generating J more often improves the mixing rate of the resulting Markov chain which offsets the minor increase in computational cost.*

3.2 Particle Marginal Metropolis Hastings

The PMMH approach in Andrieu et al. (2010) corresponds to a particular choice of proposal distribution. Andrieu et al. (2010) show that

$$\frac{\tilde{\pi}^N(U_{1:T},\theta_i|\theta_{-i})}{\psi(U_{1:T}|\theta_{-i},\theta_i)} = \frac{Z(U_{1:T},\theta)p(\theta_i|\theta_{-i})}{p(y_{1:T}|\theta_{-i})} \tag{5}$$

and hence if

$$q_{i,2}(\cdot|U_{1:T},J,\theta) = \psi(\cdot|\theta),$$

then the Metropolis-Hasting acceptance probability in (3) simplifies to

$$1 \wedge \frac{Z(\theta_i^*,\theta_{-i},U_{1:T}^*)}{Z(\theta_i,\theta_{-i},U_{1:T})} \frac{q_{i,1}(\theta_i|U_{1:T}^*,J^*,\theta_{-i},\theta_i^*)p(\theta_i^*|\theta_{-i})}{q_{i,1}(\theta_i^*|U_{1:T},J,\theta_{-i},\theta_i)p(\theta_i|\theta_{-i})}. \tag{6}$$

With this choice of Metropolis-Hasting proposal, the PMMH steps can be viewed as involving a particulate approximation to an *ideal* sampler. The approximation is effectively used to estimate the likelihood of the model. Pitt et al. (2012) show that the performance of the PMMH algorithm depends on how well the particle filter approximates the log-likelihood. They prove (under certain strong assumptions) that the optimal tradeoff between computational cost (number of particles) and mixing of the marginal chain is attained when the standard deviation of the estimated log likelihood is close to one. This version of the PMMH algorithm can also be viewed as a Metropolis-Hastings algorithm using an unbiased estimate of the likelihood.

Remark 4 *For the PMMH approach, if the proposal $q_{i,1}(\cdot|U_{1:T},J,\theta_{-i},\theta_i)$ does not involve the variable J then the acceptance probability in (3) also does not involve the variables J and J^*. In this case, it is only necessary to generate the new value J^* if the proposed value of θ_i^* is accepted.*

3.3 Updating the particles

One important application of the general algorithm is to update the particles in the PMMH algorithm using a particle Gibbs step, either every predetermined number of iterations, or with a certain probability. This is designed to prevent the algorithm getting stuck because of spurious high likelihood ratios as discussed in Pitt et al. (2012). This modification was first noticed in Andrieu et al. (2010), Section 3.1, as a way of improving the performance of the PMMH algorithm and to escape local modes. Let $\theta = (\theta_1, \ldots, \theta_p)$ be a partition of the parameter vector into p components where each component may be a vector. The mixture algorithm is defined as follows, taking $0 < p^* < 1$.

Sampling Scheme 2 *Given initial values for $U_{1:T}$, J and θ, one iteration of the MCMC involves the following steps*
1. Sample k from a Bernoulli distribution with parameter p^*
2. If $k = 0$ (PMMH sampling) For $i = 1, \ldots, p$
Step i:

(a) Sample $\theta_i^* \sim q_{i,1}(\cdot|U_{1:T}, J, \theta_{-i}, \theta_i)$.
(b) Sample $U_{1:T}^* \sim q_{i,2}(\cdot|U_{1:T}, J, \theta_{-i}, \theta_i^*)$.
(c) Sample J^* from $\tilde{\pi}^N(\cdot|U_{1:T}^*, \theta_{-i}, \theta_i^*)$.
(d) Accept the proposed values θ_i^*, $U_{1:T}^*$ and J^* with probability

$$1 \wedge \frac{\tilde{\pi}^N(U_{1:T}^*, \theta_i^*|\theta_{-i})}{\tilde{\pi}^N(U_{1:T}, \theta_i|\theta_{-i})} \frac{q_i(U_{1:T}, \theta_i|U_{1:T}^*, J^*, \theta_{-i}, \theta_i^*)}{q_i(U_{1:T}^*, \theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i)}$$

where

$$q_i(U_{1:T}^*, \theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i) = q_{i,1}(\theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i) q_{i,2}(U_{1:T}^*|U_{1:T}, J, \theta_{-i}, \theta_i^*).$$

3. else if $k = 1$ (PG sampling update)

(a) Sample $U_{1:T}^{(-J)} \sim \tilde{\pi}^N\left(\cdot|X_{1:T}^J, B_{1:T-1}^J, J, \theta\right)$ using Algorithm 2.
(b) Sample $J \sim \tilde{\pi}^N(\cdot|U_{1:T}, \theta)$.

The sequential algorithm is obtained by setting $p_1 = p$ in Sampling Scheme 1. This way, the Gibbs step of the algorithm only samples the particles and the trajectory. Other variations are possible. For example, it is possible to carry out multiple PMMH steps for every Gibbs step.

4 Examples

This section uses three example to illustrate empirically that it is important to have the flexibility of Sampling Scheme 1 given in our article. That is, it is desirable to generate parameters that are highly correlated with the states using a PMMH step that does not condition on the states. Conversely, if there is a subset of parameters that is not highly correlated with the states, then it is preferable to generate them using a particle Gibbs step, or a particle Metropolis within Gibbs step, that conditions on the states, especially when the subset is large. Although our examples are relatively simple for illustrative purposes, they capture the essence of our message and we emphasize that it is relatively straightforward to modify

the code for more complex applications that cannot be easily done with standard MCMC methods.

We use the integrated autocorrelation time (IACT) standardized for running time to compare the performance of the MCMC algorithms. In our examples this means that we choose the number of particles for each algorithm in such a way that the running times of the algorithms are about the same so that the IACTS's of the algorithms show their inefficiency, and the smaller the IACT the better. To calculate the IACT we use the method of overlapping batch means (Jones et al. 2006, Section 3.2) and report the bandwidth that we chose. We use the following notation to describe the algorithms used in the examples. The basic samplers, as used in Sampling Scheme 1, are $PMMH(\cdot)$ and $PG(\cdot)$. These samplers can be used alone or in combination. For example, $PMMH(\theta)$ means using a PMMH step to sample the parameter vector θ; $PMMH(\theta_1)+PG(\theta_2)$ means sampling θ_1 in the PMMH step and θ_2 in the PG step (with backward simulation); and $PG(\theta)$ means sampling θ using the PG sampler (with backward simulation unless stated otherwise). In all the examples, the PMMH step uses the bootstrap particle filter to sample the particles ($q_{i,2}$) and the adaptive random walk in Roberts and Rosenthal (2009) as the proposal density $q_{i,1}$ for the parameters. The particle filter is implemented in C, while the parameter samplers are implemented in Matlab.

We note that efficient Metropolis-Hastings schemes for standard MCMC are possible because both the likelihood and its derivatives can be computed exactly or, in the case of derivatives, by finite differences. This is not the case, however, for PMMH where the likelihood is estimated and its derivatives are not available even numerically. We therefore use the adaptive random walk in all our examples to generate the parameters as it provides a general and flexible sampling solution that does not require the computation or estimation of the local gradient or hessian. The random walk proposal scales the covariance matrix by a factor of $2.38/\sqrt{d}$ in the non-particle case and $2.56/\sqrt{d}$ in the particle case (Sherlock et al. 2015) , where d is the number of parameters in the random walk. This means that for large d the scale factor can be very small and the random walk may be very inefficient because it moves slowly. Moreover, and importantly, Sherlock et al. (2015) show that in the particle case, the optimal acceptance rate of random walk proposals is 7%, which suggests that it is quite inefficient. Sherlock et al. (2015) derive this acceptance rate assuming that $d \to \infty$, which also means that at least theoretically the step size is very small.

The use of more efficient particle MALA algorithms is discussed by Nemeth et al. (2014) who write "Our results show that the behaviour of particle MALA depends crucially on how accurately we can estimate the gradient of the log-posterior. If the error in the estimate of the gradient is not controlled sufficiently well as we increase dimension then asymptotically there will be no advantage in using particle MALA over a particle MCMC algorithm using a random-walk proposal." However, as noted by Sherlock et al. (2015), it is in general even more difficult to obtain accurate estimates of the gradient of the log-posterior than it is to get accurate estimates of the log-posterior.

This discussion again suggests why using particle Gibbs, or particle Metropolis within Gibbs, may be preferred to PMMH whenever possible, because it may be easier to obtain better proposals within such a particle Gibbs, or particle Metropolis within Gibbs, framework.

In Section 4.2 we also use the adaptive random walk for the non-particle ideal sampler in the Metropolis-Hastings step, even though proposals that may be more efficient than the adaptive random walk may be available here. The reason we do so is to enable a comparison between the relative performances of the PMMH, PG and PMMH+PG samplers with the relative performances of the analogous non-particle samplers.

4.1 Stochastic volatility model with covariates in the mean

We consider

$$y_t = \beta' z_t + e^{x_t/2}\,\varepsilon_t \quad \text{and} \quad x_{t+1} = \mu + \phi\,(x_t - \mu) + \eta_{t+1}, \tag{7}$$

which is a general time series model that allows for a scalar dependent variable y_t with possible dependence on covariates in the mean as well as stochastic variance terms. Thus, $E(y_t|z_t, x_t, \beta, \theta) = \beta' z_t$, where z_t can consist of lags of y_t; $\text{Var}(y_t|x_t, z_t, \beta, \theta) = \exp(x_t)$, with x_t a first order autoregression. The model can be applied to many time series and has been extensively used in the financial econometrics literature where it is known as a stochastic volatility model. It is straightforward to generalize this model in a number of ways: for example by allowing for covariates in the conditional variance and including a conditional variance term in the mean. See Durbin and Koopman (2012), pp.216-221, who discuss the basic stochastic volatility model and some extensions. Although the basic model (7) can also be estimated by regular MCMC methods, more complex versions of it that allow for jumps in the volatility will require particle methods. See, for example, the model in eq. (4) in Pitt et al. (2014). However, the smooth particle filter used by these authors does not have the flexibility of Sampling Scheme 1 of generating some of the parameters conditional on the states.

In our article, for purposes of illustration, we confine the discussion to the basic model given by (7), where $\varepsilon_t \sim N(0,1)$, $\eta_t \sim N(0,\tau^2)$ and $x_1 \sim N(\mu, \tau^2/(1-\phi^2))$. The notation $N(a,b)$ means a normal distribution with mean a and variance b. The unknown parameters in the model are β and $\theta = (\tau, \mu, \phi)$, with τ, μ and ϕ scalar and β an $m \times 1$ vector.

We compare the $PMMH(\theta, \beta)$, the $PG(\theta, \beta)$ and the $PMMH(\theta) + PG(\beta)$ samplers and show that the best results are obtained by generating the state evolution density parameters μ, ϕ, and τ not conditioning on the states and β conditional on the states.

The dimension of the covariate vector is $m = 50$ and the covariates are drawn from a multivariate Gaussian distribution $z_t \sim N(0, .5^2 I_{50})$. We simulate $T = 500$ observations from (7) with parameters $\mu = -.2$, $\phi = .98$, $\tau = .15$, and $\beta_1 = \cdots = \beta_n = 1$. The n-dimensional vector β is assumed to have a diffuse prior, i.e., the prior is $p(\beta) \propto 1$ and the priors for the state evolution density are $\mu \sim N(0, 100)$, $\phi \sim U(-1, 1)$, and $\tau \sim IG(2.5, 0.05)$. We run all sampling schemes for 50,000 iterations and discard the initial 20,000 as warmup. All algorithms use the starting values $(\mu, \phi, \tau) = (0, 0, .5)$ and $\beta_1 = \cdots = \beta_m = 0$.

Table 1 summarizes the results of the simulation. We choose the number of particles for each sampler so that the computing times are about the same, and hence the efficiency of the samplers can be compared in terms of their IACT values. The table shows that the $PMMH(\theta, \beta)$ mixes the worst of all three sampling schemes and its acceptance rate is very low. The $PG(\theta, \beta)$ sampler performs well for β, being almost as good as an independent sampler. However, the standard deviation τ of the disturbance η_{t+1} in the state transition equation has a large IACT of 90. The $PMMH(\theta) + PG(\beta)$ sampler performs as well as the PG sampler for β, but its IACT for τ is much lower.

4.2 Semiparametric model: Management of a retirement fund

This example use a linear Gaussian state space model to compare several particle MCMC samplers with their ideal counterparts. By an ideal sampler we mean one where the MCMC is carried out without using an particle approximations. This example shows that the relative performance of the ideal samplers also holds for their particle counterparts, and in particular

	$PMMH(\theta,\beta)$	$PG(\theta,\beta)$	$PMMH(\theta)+PG(\beta)$
$IACT(\mu)$	>232	1.91	42.98
$IACT(\phi)$	>368	34.36	28.24
$IACT(\tau)$	>368	89.65	24.58
$\min_{i=1:50} IACT(\beta_i)$	>262	1.00	1.00
$\max_{i=1:50} IACT(\beta_i)$	>375	1.91	2.10
time/1000 iterations	79.22	79.95	76.76
# particles	460	380	200
acc. rate	3.5%	–	14.8%

Table 1: Summary of the simulation results for the stochastic volatility model with regressors. The table reports: (i) the IACT estimates for μ,ϕ,τ, and the minimum and maximum estimated IACT's for the β_i's; (ii) time to run 1000 iterations; (iii) number of particles; and (iv) acceptance rates. The IACT can be directly compared as the time to run each algorithm is approximately the same. The IACT's are estimated using the method of overlapping batch means with a bandwidth of 500 samples. This bandwidth seems to underestimate the IACT for the PMMH algorithm and is shown only for reference.

we show that it is desirable to use Sampling Scheme 1. Although we have chosen to use an example with Gaussian errors to make it possible to compare the particle and ideal samplers, it is straightforward to use the same model for the mean with a general link function, for example, the binomial model in Section 4.3, where it is advantageous to use particle methods.

We consider the retirement fund data, used and discussed in more detail in Ruppert et al. (2003), Section 1.3, and available at http://matt-wand.utsacademics.info/webspr/data.html. It concerns a company named Best Retirement Inc. (BRI) that sells retirement plans to corporations around the United States. The dataset relates to 401(k) retirement plans for employees of 92 firms, and data is available on several attributes of the firms from previous years. BRI's main goal is to estimate the year-end dollar amount contributed to each plan in advance so that it can make internal revenue and cost projections.

The dependent variable is the contribution to the retirement plan at the end of the first year (*contrib*) using as explanatory variables whether the client has group life or health insurance (*group*), employee turnover rate (*turnover*), number of employees eligible to participate in 401(k) plans (*eligible*), whether the plan has immediate vesting of employer contribution (*vest*), whether the plan has a failsafe provision (*failsafe*), percentage of provision matched by employer (*match*), annual average employee salary (*salary*), underwriter's estimate of end-of-year contribution in dollars (*estimate*), and whether the plan was sold by Susan, a specialist in 401(k) plans, or not (*susan*). Ruppert et al. (2003), Sections 1.3 and 18.4, note that the response variable *contrib* is heavily skewed and use its logarithm as the dependent variable. The effect of *salary* on $log(contrib)$ is nonlinear and it is estimated nonparametrically using a cubic smoothing spline.

We follow Ruppert et al. (2003) and use the semiparametric model

$$\log(contrib) = \beta' z_t + g(s_t) + e_t, \tag{8}$$

where $\beta = (\beta_1,\ldots,\beta_8)'$, z_t is an 8×1 vector of the covariates listed above and in that order, and s_t is the average salary of a firm rescaled to be between 0 and 1, i.e.,

$$s_t = salary_t/(\max\{salary_k.k = 1,\ldots,T\}\ .$$

The data is reordered to be in ascending order of s_t. The function $g(\cdot)$ is assumed to be smooth, but otherwise unknown. We place a smoothness prior on g which is given by equa-

tions (18) and (19) in Appendix C. The smoothness prior leads to the posterior mean estimate of g being a cubic smoothing spline (see Appendix C and Wahba (1978)). From Appendix C, we can write (8) in state space form with state vector $X_t = (g(s_t), dg(s_t)/ds)'$, $g(s_t) = \begin{pmatrix}1 & 0\end{pmatrix} X_t$, state transition equation (20), and with diffuse initial conditions $p(X_1) \propto 1$. The prior for β is diffuse, i.e., $p(\beta) \propto 1$, the prior for σ^2 is $IG(1,1)$, and the prior for τ^2 is $IG(1,1)$ where τ^2 is the scale parameter in the state transition equation (20).

We estimated the model using the $PG(\sigma,\tau,\beta,x_1)$, $PMMH(\sigma,\tau,\beta,x_1)$ and the $PMMH(\tau,x_1)+PG(\sigma,\beta)$ samplers. We sample $\log\tau$ and $\log\sigma$ in the PMMH step, running 50,000 iterations and discarding the initial 10,000 as warmup. We also estimated the model by their non-particle MCMC counterparts, which we call the ideal samplers, and where we denote the Gibbs sampler by $GS(\sigma,\tau,\beta,x_1)$, the random walk Metropolis sampler by $RWM(\sigma,\tau,\beta,x_1)$, and the hybrid sampler which combines a random walk Metropolis step with a Gibbs sampling step by $RWM(\tau,x_1)+GS(\sigma,\beta)$. As stated above, the aims of this example are to show that the relative performance of the ideal samplers also holds for their particle counterparts, and in particular we show the desirability of using the hybrid Sampling Scheme 1.

Table 2 summarizes the simulation results. The table shows that $PMMH(\tau,x_1)+PG(\sigma,\beta)$ outperforms $PMMH(\sigma,\tau,\beta,x_1)$ because its IACT's are lower, and confirms the desirability of using Sampling Scheme 1.

All the IACT estimates are obtained using the method of overlapping batch means, with window size 200. The number of particles for each sampler is chosen so that the running times of the samplers are about the same. We do not report the results for the particle Gibbs sampler because it did not converge in 50,000 iterations. The bad performance of the particle Gibbs algorithm is not unexpected as it well established in the literature that the full Gibbs sampler, that samples τ conditional on the states, performs poorly (Carter and Kohn 1994, 1996).

	$PMMH(\sigma,\tau,\beta,x_1)$	$PMMH(\tau,x_1)+PG(\sigma,\beta)$
$IACT(\widehat{\sigma})$	51.16	2.37
$IACT(\widehat{\tau})$	54.36	29.26
$\min_{i=1:8} IACT(\widehat{\beta_i})$	44.30	1.00
$\max_{i=1:8} IACT(\widehat{\beta_i})$	52.30	3.71
Time/1000 iterations	14.44	15.08
# particles	400	200
acc. rate	20.83%	18.70%

Table 2: Results for semiparametric Model: Performance of the $PMMH(\sigma,\tau,\beta,x_1)$ and $PMMH(\tau,x_1)+PG(\sigma,x_1)$ samplers. The table shows: (i) the IACT for σ, τ, and the minimum and maximum IACT for the β's; (ii) the actual running time in seconds per 1000 iterations; and (iii) the number of particles and acceptance rate of the PMMH step for each sampler.

The posterior means and standard deviations (in brackets) of the parameters, estimated using the $PMMH(\tau,x_0)+PG(\sigma,\beta)$ sampler are: $\widehat{\sigma} = 0.7033\,(.0004)$, $\widehat{\tau} = 1.9328\,(.0884)$, $\widehat{\beta_1} = -.1199\,(.0011)$, $\widehat{\beta_2} = -.0867\,(.0008)$, $\widehat{\beta_3} = .0613\,(.0027)$, $\widehat{\beta_4} = -.0279\,(.0008)$, $\widehat{\beta_5} = -.0168\,(.0009)$, $\widehat{\beta_6} = -.0481\,(.0010)$, $\widehat{\beta_7} = .6110\,(.0024)$, $\widehat{\beta_8} = .1115\,(.0008)$, and the initial value x_0 is estimated as $\widehat{x}_{1,1} = 13.1341\,(0.0067)$ and $\widehat{x}_{1,2} = 0.4487\,(0.0205)$. These values are consistent across all samplers.

Table 3 summarizes the results for the corresponding ideal samplers $RWM(\sigma,\tau,\beta,x_1)$ and $RWM(\tau,x_1)+GS(\sigma,\beta)$ and shows that they have similar inefficiency factors to the particle samplers.

	$RWM(\sigma,\tau,\beta,x_1)$	$RWM(\tau,x_1)+GS(\sigma,\beta)$
$IACT(\widehat{\sigma})$	44.24	1.85
$IACT(\widehat{\tau})$	53.14	26.84
$\min_{i=1:8} IACT(\widehat{\beta}_i)$	35.10	1.00
$\max_{i=1:8} IACT(\widehat{\beta}_i)$	45.18	2.93
acc. rate	22.90%	21.03%

Table 3: Semiparametric Model: Performance of the $RWM(\sigma,\tau,\beta,x_1)$ and $RWM(\tau,x_1)+GS(\tau,x_1)$ samplers. The table shows: (i) the IACT for σ, τ, and the minimum and maximum IACT for the β's; (ii) the actual running time in seconds per 1000 iterations; and (iii) the number of particles and acceptance rate of the PMMH step for each sampler.

4.3 Binomial regression with shifts in the mean

There is a great deal of interest in modeling time series subject to interventions such as structural breaks and regime switching. For example, see early work on change-point problems by Barry and Hartigan (1993), Kim and Nelson (1999) on regime switching models, and Erdman and Emerson (2008) for an application to microarray data. Although standard MCMC methods can deal with change point problems for conditionally Gaussian time series models, e.g. Gerlach et al. (2000) and Giordani et al. (2007), it is much more challenging to use them for general time series regression problems, because it is hard, in general, to integrate out the states. The purpose of our third example is to show that particle methods offer an attractive alternative for dealing with change point problems in general time series models because it is easy to integrate out the states.

To illustrate, we consider the binomial time series regression with logit link

$$\begin{aligned} y_t|(n_t,x_t,z_t,\theta) &\sim \text{Bin}(n_t,p_t), \\ \text{with} \quad p_t &= \frac{\exp(x_t+\beta' z_t)}{1+\exp(x_t+\beta' z_t)}, \\ \text{i.e.,} \quad \text{logit}(p_t) &= x_t+\beta' z_t \end{aligned}$$

where we allow the intercept x_t to change over time, through the state transition equation

$$x_{t+1} = x_t + k_{t+1}\eta_{t+1},$$

where the η_t are independent $N(0,\tau^2)$ and $k_t, t \geq 2$, is a binary Markov process. We discuss two scenarios: (a) In the first scenario $k_t = 0$ or 1 with $P(k_t = 1) = \omega$. Then the prior for x_t allows for a piecewise constant process with transitions whose number and location is unknown. (b) In the second scenario, $k_t = 1$ or 100, with $P(k_t = 100) = \omega$. This model allows for a continuous evolution of x_t if $k_t = 1$ or allows a jump in x_t if $k_t = 100$. In both scenarios, it is cumbersome to carry out a standard MCMC unless the n_t are small, even if we use latent variable methods such as those in Polson et al. (2013) or Frühwirth-Schnatter and Frühwirth

(2007). For Scenario (b), we can use the particle filter to integrate out both the k_t and the x_t, which is more efficient than the approach in Gerlach et al. (2000) who, in the context of conditionally Gaussian models, integrate out the x_t and generate the k_t one at a time. We can apply sampling Scheme 1 to both Scenarios (a) and (b), but in Scenario (a) it is now necessary to use a mixed measure rather than the Lebesgue measure used in our derivations. It is straightforward to extend our results to the mixed measure case.

We now study numerically the relative performance of Sampling scheme 1 under Scenario (a) with the k_t independent. However, the approach extends in a straightforward way to time series models with a general observation equation $p(y_t|x_t)$ which we assume can be evaluated given x_t, as well as allowing the k_t to be a binary Markov process, where the presence of a change point at time t affects the probability of a change point at time $t+1$.

In our example we treat x_1, the initial value of x_t, as an unknown parameter with the fairly noninformative prior $x_1 \sim N(0,10)$. We denote the vector of the other parameters as $\theta = (\tau^2, \omega, \beta)$ and take their prior distributions as $\tau^2 \sim IG(2,5), \omega \sim \text{Beta}(8,2)$ and $p(\beta) \propto 1$ and independent of each other and x_1. The prior for ω means that the median is close to 0.82 and 90% of its probability mass is greater than 0.63.

We simulate $T = 300$ observation from the true model with $n_t \sim \text{Bin}(10,.5)$, $\tau^2 = 5$, $\omega = .87$, and $\beta = (-1.5, 1.0)'$. The covariates are drawn independently from a standard normal distribution. We compare the performance of the estimates using three samplers: PMMH, a particle Gibbs sampler with ancestral tracing, and a hybrid sampler of the form $PG(\beta) + PMMH(\tau^2, \omega)$. We use $N = 400$ particles in the first two samplers and $N = 200$ in the hybrid sampler. We ran all samplers for 50,000 iterations and discarded the initial 10,000 as warmup. We started the chains at the true parameter values.

Table 4 summarizes the results for the three samplers and shows that both $PMMH$ and $PMMH + PG$ samplers perform better than the PG sampler. The $PMMH + PG$ sampler performs slightly better than the PMMH sampler even when normalized by time. The IACT estimates were calculated using overlapping batch means with bandwidth of 578 samples.

	$PG(\theta)$	$PMMH(\theta)$	$PMMH(\tau^2,\omega)+PG(\beta)$
IACT(τ^2)	106.4	25.8	18.8
IACT(ω)	166.0	44.7	33.5
IACT(β_1)	28.3	31.7	31.2
IACT(β_2)	24.6	32.2	25.5
Time/1000 iterations	19.4	17.4	19.1
# particles	400	400	200
acc. rate	–	22.6%	15.7%

Table 4: Bimomial model with two regressors. Performance of the $PMMH(\theta)$, $PG(\theta)$, and $PMMH(\tau^2,\omega)+PG(\beta)$ samplers. The table shows the IACT for all the parameters in the model. The actual running time in seconds per 1000 iterations, the number of particles and acceptance rate of the PMMH step are also given for each algorithm.

As we increase the number of regressors, the performance of the $PMMH + PG$ sampler improves relative to the PMMH sampler. Table 5 shows the simulation results for $m = 4$ and $m = 8$ regressors. The regression coefficients are $\beta = (-1.5, 0, -1.5, 0)$ for $m = 4$ and $\beta = (-1.5, 0, -1.5, 0, -1.5, 0, -1.5, 0)$ for $m = 8$; the τ^2 and ω parameters are the same as for the $m = 2$ case. For a given number of regressors, the inefficiency factors for each parameter are directly comparable across samplers. The table shows that the performance of the $PMMH$

sampler degenerates as the number of parameters increases, while the $PMMH+PG$ is more robust to this change.

m		$PMMH(\theta)$	$PMMH(\tau^2,\omega)+PG(\beta)$
4	IACT(τ^2)	47.5	24.8
	IACT(ω)	81.1	41.9
	$\max_{1\leq i\leq m} IACT(\beta_i)$	89.7	46.8
	Time/1000 iterations	17.6	20.8
	# particles	400	200
	acc. rate	21.6%	15.9%
8	IACT(τ^2)	47.4	13.4
	IACT(ω)	147.2	28.8
	$\max_{1\leq i\leq m} IACT(\beta_i)$	94.0	62.4
	Time/1000 iterations	17.6	23.8
	# particles	400	200
	acc. rate	21.7%	20.0%

Table 5: Binomial example for $m=4$ and $m=8$ regressors. Performance of the $PMMH(\theta)$ and $PMMH(\tau^2,\omega)+PG(\beta)$ samplers. $PG(\theta)$ is not included because it performed badly for (τ^2,ω) compared to the other two methods. The table shows: (i) the IACT estimates for τ^2, ω, and the largest IACT among β_i's ($i=1,\ldots,p$); (ii) the actual running time in seconds per 1000 iterations; (iii) the number of particles; and (iv) the acceptance rate of the PMMH steps for each sampler.

5 Discussion

Our article introduces a flexible particle Markov chain Monte Carlo sampling scheme for state space models where some parameters are generated without conditioning on the states while other parameters are generated conditional on the states. Previous schemes generated either one way or the other without combining both strategies. The technical contribution of our article is to set out the required particle framework for the general sampler and to obtain uniform ergodicity under given assumptions. Our examples demonstrate that it is advantageous to use this flexible sampling scheme to generate the parameters that are highly correlated with the states without conditioning on the states (the PMMH component) while other parameters are generated by particle Gibbs (PG) or particle Metropolis within Gibbs (PMwG).

Acknowledgement

We would like to thank the three anonymous reviewers who helped to improve the presentation and technical content of the paper. The work of the authors was partially supported by an ARC Research Council Grant DP120104014.

References

C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009.

C. Andrieu and M. Vihola. Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. arXiv preprint arXiv:1210.1484, 2012.

C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statiical Society, Series B*, 72:269–342, 2010.

D. Barry and J. A. Hartigan. A Bayesian analysis for change point problems. *Journal of the American Statistical Association*, 88(421):309–319, 1993.

B. P. Carlin, N. G. Polson, and D. S. Stoffer. A Monte Carlo approach to nonnormal and nonlinear state-space modeling. *Journal of the American Statistical Association*, 87(418):493–500, 1992.

C. Carter and R. Kohn. Markov chain Monte Carlo in conditionally Gaussian state space models. *Biometrika*, 83(3):589–601, 1996.

C. K. Carter and R. Kohn. On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553, 1994.

N. Chopin and S. S. Singh. On the particle Gibbs sampler. arXiv preprint arXiv:1304.1887, 2013.

R. Douc and O. Cappé. Comparison of resampling schemes for particle filtering. In *Image and Signal Processing and Analysis, 2005. ISPA 2005. Proceedings of the 4th International Symposium on*, pages 64–69. IEEE, 2005.

A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000.

A. Doucet, M. Pitt, G. Deligiannidis, and R. Kohn. Efficient implementation of markov chain monte carlo when using an unbiased likelihood estimator. *Biometrika*, 2015. forthcoming.

J. Durbin and S. Koopman. *Time series analysis of state space methods*. Oxford University Press, second edition, 2012.

C. Erdman and J. Emerson. A fast Bayesian change point analysis for the segmentation of microarray data. *Bioinformatics*, 24(19):2143–48, 2008.

S. Frühwirth-Schnatter. Data augmentation and dynamic linear models. *Journal of Time Series Analysis*, 15 (2):183–202, 1994.

S. Frühwirth-Schnatter. *Finite Mixture and Markov Switching Models*. New York, NY: Springer Science+ Business Media, LLC, 2006.

S. Frühwirth-Schnatter and R. Frühwirth. Auxiliary mixture sampling with applications to logistic models. *Computational Statistics and Data Analysis*, 51:3509–3528, 2007.

R. Gerlach, C. Carter, and R. Kohn. Efficient Bayesian inference for dynamic mixture models. *Journal of the American Statistical Association*, 95(451):819–828, 2000.

P. Giordani, R. Kohn, and D. van Dijk. A unified approach to nonlinearity, structural change, and outliers. *Journal of Econometrics*, 137(1):112–133, 2007.

S. Godsill, A. Doucet, and M. West. Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association*, 99(465):156–168, 2004.

N. J. Gordon, D. J. Salmond, and A. F. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE Proceedings F (Radar and Signal Processing)*, volume 140, pages 107–113. IET, 1993.

D. Guo, X. Wang, and R. Chen. New sequential Monte Carlo methods for nonlinear dynamic systems. *Statistics and computing*, 15:135–147, 2005.

G. L. Jones, M. Haran, B. S. Caffo, and R. Neath. Fixed-width output analysis for Markov chain Monte Carlo. *Journal of the American Statistical Association*, 101(476):1537–1547, 2006.

C.-J. Kim and C. R. Nelson. *State-space models with regime switching: classical and Gibbs-sampling approaches with applications*, volume 1 of *MIT Press Books*. The MIT press, 1999.

S. Kim, N. Shephard, and S. Chib. Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3):361–393, 1998.

G. Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of computational and graphical statistics*, 5(1):1–25, 1996.

F. Lindsten and T. B. Schön. On the use of backward simulation in particle Markov chain Monte Carlo methods. arxiv:1110.2873, 2012.

F. Lindsten, M. I. Jordan, and T. B. Schön. Particle Gibbs with ancestor sampling. arxiv:1401.0604, 2014.

J. Liu and M. West. Combined parameter and state estimation in simulation-based filtering. In N. de Freitas, A. Doucet, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*, pages 197 – 223. Springer-verlag, New York, 2001.

C. Nemeth, C. Sherlock, and P. Fearnhead. Particle metropolis adjusted langevin algorithms, 2014. http://arxiv.org/abs/1412.7299.

J. Olsson and T. Ryden. Rao-Blackwellization of particle Markov chain Monte Carlo methods using forward filtering backward sampling. *Signal Processing, IEEE Transactions on*, 59(10):4606–4619, 2011.

M. Pitt, S. Malik, and A. Doucet. Simulated likelihood inference for stochastic volatility models using continuous particle filtering. *Annals of Institute of Statistical Mathematics*, 66(421):527–552, 2014.

M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the American statistical association*, 94(446):590–599, 1999.
M. K. Pitt, R. d. S. Silva, P. Giordani, and R. Kohn. On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics*, 171:134–151, 2012.
N. G. Polson, J. R. Stroud, and P. Müller. Practical filtering with sequential parameter learning. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(2):413–428, 2008.
N. G. Polson, J. G. Scott, and J. Windle. Bayesian inference for logistic models using pólya–gamma latent variables. *Journal of the American Statistical Association*, 108(504):1339–1349, 2013.
G. O. Roberts and J. S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1:20–71, 2004.
G. O. Roberts and J. S. Rosenthal. Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18(2):349–367, 2009.
D. Ruppert, M. Wand, and R. Carroll. *Semiparametric regression*. Cambridge University Press, 2003.
N. Shephard and M. K. Pitt. Likelihood analysis of non-Gaussian measurement time series. *Biometrika*, 84 (3):653–667, 1997.
C. Sherlock, A. Thiery, G. Roberts, and J. Rosenthal. On the efficieny of pseudo-marginal random walk metropolis algorithms. *Annals of Statistics*, 2015. forthcoming.
G. Storvik. Particle filters for state-space models with the presence of unknown static parameters. *Signal Processing, IEEE Transactions on*, 50(2):281–289, 2002.
L. Tierney. A note on Metropolis-Hastings kernels for general state spaces. *The Annals of Applied Probability*, 8(1):1–9, 1998.
R. Van Der Merwe, A. Doucet, N. De Freitas, and E. Wan. The unscented particle filter. *Advances in neural information processing systems*, pages 584–590, 2001.
D. van Dyk and T. Park. Partially collapsed Gibbs samplers: theory and methods. *Journal of the American Statistical Association*, 103:790–796, 2008.
G. Wahba. Improper priors, spline smoothing and the problem of guarding against model errors in regression. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 364–372, 1978.
W. E. Wecker and C. F. Ansley. The signal extraction approach to nonlinear regression and spline smoothing. *Journal of the American Statistical Association*, 78(381):81–89, 1983.

A Ergodicity

This appendix first discusses convergence in total variation norm and then considers the stronger condition of uniform convergence.

Note that, by construction, Sampling Scheme 1 has the stationary distribution

$$\tilde{\pi}^N\left(x_{1:T}, b_{1:T-1}, j, u_{1:T}^{(-j)}, \theta\right)$$

defined in (2). From Roberts and Rosenthal (2004) Theorem 4, irreducibility and aperiodicity are sufficient conditions for the Markov chain obtained using Sampling Scheme 1 to converge to its stationary distribution in total variation norm for $\tilde{\pi}^N$-almost all starting values. These conditions must be checked for a particular sampler and it is often straightforward to do so. We will relate Sampling Scheme 1 to the particle Metropolis within Gibbs sampling scheme defined below.

Sampling Scheme 3 (Ideal) *Given initial values for $U_{1:T}$, J and θ, one iteration of the MCMC involves the following steps*
1. (PMMH sampling) For $i = 1, \ldots, p_1$
Step i:

(a) Sample $\theta_i^* \sim q_{i,1}(\cdot|U_{1:T}, J, \theta_{-i}, \theta_i)$.
(b) Sample $\left(J^*, U_{1:T}^*\right) \sim \tilde{\pi}^N(\cdot|\theta_{-i}, \theta_i^*)$.
(c) Accept the proposed values θ_i^*, $U_{1:T}^*$ and J^* with probability

$$\widetilde{\alpha}_i\left(U_{1:T}, J, \theta_i; U_{1:T}^*, J^*, \theta_i^* | \theta_{-i}\right) = 1 \wedge \frac{\tilde{\pi}^N(\theta_i^*|\theta_{-i})}{\tilde{\pi}^N(\theta_i|\theta_{-i})} \frac{q_{i,1}(\theta_i|U_{1:T}^*, J^*, \theta_{-i}, \theta_i^*)}{q_{i,1}(\theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i)} \tag{9}$$

2. (PG or PMwG sampling) For $i = p_1 + 1, \ldots, p$
Step i:

(a) Sample $\theta_i^* \sim q_i(\cdot|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i)$.
(b) Accept the proposed value θ_i^* *with probability*

$$\begin{aligned}&\alpha_i\left\{\theta_i;\theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right\} = \\ &1 \wedge \frac{\tilde{\pi}^N\left(\theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right)}{\tilde{\pi}^N\left(\theta_i|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right)} \\ &\frac{q_i(\theta_i|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i^*)}{q_i(\theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i)}.\end{aligned} \tag{10}$$

3. *Sample* $U_{1:T}^{(-J)} \sim \tilde{\pi}^N(\cdot|X_{1:T}^J, B_{1:T-1}^J, J, \theta)$ *using Algorithm 2.*
4. *Sample* $J \sim \tilde{\pi}^N(\cdot|U_{1:T}, \theta)$.

We call Sampling Scheme 3 an *ideal* particle sampling scheme because in Part 1 Step *i*(b) it generates the particles $U_{1:T}^*$ from their conditional distribution $\tilde{\pi}^N(\cdot|\theta_{-i}, \theta_i^*)$ instead of using a Metropolis-Hasting proposal. Thus comparing Sampling Schemes 1 and 3 allows us to concentrate on the effect of the Metropolis-Hastings proposal for the particles on the convergence of the sampler.

Remark 5 *Andrieu and Roberts (2009) and Andrieu and Vihola (2012) discuss the relationship between PMMH sampling schemes with one block of parameters and an ideal Metropolis-Hastings sampling scheme not involving the particles. Sampling Schemes 1 and 3 are more general. Our approach, however, is similar to and generalizes the results in Andrieu and Roberts (2009) and Andrieu and Vihola (2012) to more complex sampling schemes.*

To develop the theory of Sampling Schemes 1 and 3 we require the following definitions. Let $\left\{V^{(n)}, n = 1, 2, \ldots\right\}$ be the iterates of a Markov chain defined on the state space $\mathscr{V} := \mathscr{U} \times \mathbb{N} \times \Theta$.

Let $K(v;\cdot)$ be the substochastic transition kernel of Sampling Scheme 1 that defines the probabilities for accepted Metropolis-Hastings moves. Note that probabilities involving the substochastic kernel provide lower bounds on the probabilities for the transition kernel of the corresponding Markov chain. Note also that K it is the product of substochastic transition kernels for each of the p steps of the sampling scheme, which we denote by

$$K = K_1 K_2 \ldots K_p.$$

Note that for $i = 1, \ldots, p_1$

$$\begin{aligned}&K_i\left(U_{1:T}, J, \theta_{-i}, \theta_i; U_{1:T}^*, J_i^*, \theta_{-i}, \theta_i^*\right) = \\ &\tilde{\pi}^N(J^*|U_{1:T}^*, \theta_{-i}, \theta_i^*) q_i(U_{1:T}^*, \theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i) \times \\ &\alpha_i\left(U_{1:T}, J, \theta_i; U_{1:T}^*, J^*, \theta_i^*|\theta_{-i}\right)\end{aligned}$$

Similarly, let $\widetilde{K}(v;\cdot)$ be the substochastic transition kernel of Sampling Scheme 3 that define the probabilities for accepted Metropolis-Hastings moves. It is also the product of substochastic transition kernels for each step of the ideal sampling scheme denoted by

$$\widetilde{K} = \widetilde{K}_1 \widetilde{K}_2 \ldots \widetilde{K}_p.$$

Note that the kernels K_i and $\widetilde{K}_i$ only differ for $i = 1, \ldots, p_1$.

A sufficient condition for Sampling Scheme 1 to be irreducible and aperiodic is given in the following theorem.

Theorem 2 *Suppose that* $\widetilde{K}$ *is irreducible and aperiodic and that for* $i = 1, \ldots, p_1$ *and for all* $(U_{1:T}, J, \theta_{-i}, \theta_i^*) \in \mathscr{V}$

$$\tilde{\pi}^N(\cdot|\theta_{-i}, \theta_i^*) \ll q_{i,2}(\cdot|U_{1:T}, J, \theta_{-i}, \theta_i^*). \tag{11}$$

Then K *is irreducible and aperiodic.*

Proof. From (11), $\widetilde{K}_i \ll K_i$ for $i = 1, \ldots, p_1$ and hence $\widetilde{K} \ll K$. So $\widetilde{K}$ irreducible and aperiodic implies that K is irreducible and aperiodic. ■

Theorem 2 gives the following corollary, which is similar to Theorem 1 of Andrieu and Roberts (2009).

Corollary 3 *Suppose that $\widetilde{K}$ is irreducible and aperiodic and that for $i = 1, \ldots, p_1$,*

$$q_{i,2}(\cdot|U_{1:T}, J, \theta_{-i}, \theta_i^*) = \psi(\cdot|\theta_{-i}, \theta_i^*).$$

Then K is irreducible and aperiodic.

Proof. This follows from Lemma 1 and Theorem 2. ■

We now consider uniform ergodicity of the sampling schemes. We follow the approach in Andrieu and Roberts (2009) and give sufficient conditions for the existence of minorization conditions for Sampling Scheme 1. From Roberts and Rosenthal (2004) Theorem 8, these minorization conditions are equivalent to uniform ergodicity. The results use the following technical lemmas.

Lemma 4 *For $i = 1, \ldots, p_1$,*

$$\begin{aligned}
&\alpha_i(U_{1:T}, J, \theta_i; U_{1:T}^*, J^*, \theta_i^*|\theta_{-i}) \geq \\
&\left\{1 \wedge \frac{\tilde{\pi}^N(U_{1:T}^*|\theta_{-i}, \theta_i^*)\, q_{i,2}(U_{1:T}|U_{1:T}^*, J^*, \theta_{-i}, \theta_i)}{\tilde{\pi}^N(U_{1:T}|\theta_{-i}, \theta_i)\, q_{i,2}(U_{1:T}^*|U_{1:T}, J, \theta_{-i}, \theta_i^*)}\right\} \times \\
&\widetilde{\alpha}_i(U_{1:T}, J, \theta_i; U_{1:T}^*, J^*, \theta_i^*|\theta_{-i})
\end{aligned}$$

Proof. From (3)

$$\begin{aligned}
&\alpha_i(U_{1:T}, J, \theta_i; U_{1:T}^*, J^*, \theta_i^*|\theta_{-i}) \\
&= 1 \wedge \frac{\tilde{\pi}^N(U_{1:T}^*, \theta_i^*|\theta_{-i})}{\tilde{\pi}^N(U_{1:T}, \theta_i|\theta_{-i})} \frac{q_i(U_{1:T}, \theta_i|U_{1:T}^*, J^*, \theta_{-i}, \theta_i^*)}{q_i(U_{1:T}^*, \theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i)} \\
&= 1 \wedge \frac{\tilde{\pi}^N(U_{1:T}^*|\theta_{-i}, \theta_i^*)\, q_{i,2}(U_{1:T}|U_{1:T}^*, J^*, \theta_{-i}, \theta_i)}{\tilde{\pi}^N(U_{1:T}|\theta_{-i}, \theta_i)\, q_{i,2}(U_{1:T}^*|U_{1:T}, J, \theta_{-i}, \theta^*)} \times \\
&\quad \frac{\tilde{\pi}^N(\theta_i^*|\theta_{-i})\, q_{i,1}(\theta_i|U_{1:T}^*, J^*, \theta_{-i}, \theta_i^*)}{\tilde{\pi}^N(\theta_i|\theta_{-i})\, q_{i,1}(\theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i)} \\
&\geq \left\{1 \wedge \frac{\tilde{\pi}^N(U_{1:T}^*|\theta_{-i}, \theta_i^*)}{\tilde{\pi}^N(U_{1:T}|\theta_{-i}, \theta_i)} \frac{q_{i,2}(U_{1:T}|U_{1:T}^*, J^*, \theta_{-i}, \theta_i)}{q_{i,2}(U_{1:T}^*|U_{1:T}, J, \theta_{-i}, \theta^*)}\right\} \times \\
&\quad \left\{1 \wedge \frac{\tilde{\pi}^N(\theta_i^*|\theta_{-i})\, q_{i,1}(\theta_i|U_{1:T}^*, J^*, \theta_{-i}, \theta_i^*)}{\tilde{\pi}^N(\theta_i|\theta_{-i})\, q_{i,1}(\theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i)}\right\} \\
&= \left\{1 \wedge \frac{\tilde{\pi}^N(U_{1:T}^*|\theta_{-i}, \theta_i^*)\, q_{i,2}(U_{1:T}|U_{1:T}^*, J^*, \theta_{-i}, \theta_i)}{\tilde{\pi}^N(U_{1:T}|\theta_{-i}, \theta_i)\, q_{i,2}(U_{1:T}^*|U_{1:T}, J, \theta_{-i}, \theta_i^*)}\right\} \times \\
&\quad \widetilde{\alpha}_i(U_{1:T}, J, \theta_i; U_{1:T}^*, J^*, \theta_i^*|\theta_{-i})
\end{aligned}$$

as required. ■

Lemma 5 *Suppose that for $i = 1, \ldots, p_1$*

$$\frac{\tilde{\pi}^N(U_{1:T}^*|\theta)}{q_{i,2}(U_{1:T}^*|U_{1:T}, J, \theta)} \leq \gamma_i < \infty \tag{12}$$

for all $U_{1:T}^* \in \mathscr{U}, (U_{1:T}, J, \theta) \in \mathscr{V}$. Then, for $i = 1, \ldots, p_1$, each Markov transition kernel K_i satisfies

$$K_i \geq \gamma_i^{-1} \widetilde{K}_i \tag{13}$$

and hence

$$K \geq \left(\prod_{i=1}^{p_1} \gamma_i\right)^{-1} \widetilde{K}. \tag{14}$$

Proof. Fix $i \in \{1,\dots,p_1\}$ and let $A \in \mathscr{B}(\mathscr{U})$, $J, J^* \in \{1,\dots,N\}$ and $B \in \mathscr{B}(\Theta_i)$. Then

$$\begin{aligned}
&K_i\left(U_{1:T}, J, \theta_{-i}, \theta_i; A, J^*, \theta_{-i}, B\right) \\
&= \int_{A\times B} \tilde{\pi}^N(J^*|U^*_{1:T}, \theta_{-i}, \theta^*_i) q_i(U^*_{1:T}, \theta^*_i|U_{1:T}, J, \theta_{-i}, \theta_i) \times \\
&\quad \alpha_i\left(U_{1:T}, J, \theta_i; U^*_{1:T}, J^*, \theta^*_i|\theta_{-i}\right) \mathrm{d}U^*_{1:T} d\theta^*_i \\
&\geq \int_{A\times B} \tilde{\pi}^N(J^*|U^*_{1:T}, \theta_{-i}, \theta^*_i) q_i(U^*_{1:T}, \theta^*_i|U_{1:T}, J, \theta_{-i}, \theta_i) \times \\
&\quad \left\{1 \wedge \frac{\tilde{\pi}^N\left(U^*_{1:T}|\theta_{-i}, \theta^*_i\right) q_{i,2}(U_{1:T}|U^*_{1:T}, J^*, \theta_{-i}, \theta_i)}{\tilde{\pi}^N\left(U_{1:T}|\theta_{-i}, \theta_i\right) q_{i,2}(U^*_{1:T}|U_{1:T}, J, \theta_{-i}, \theta^*_i)}\right\} \times \\
&\quad \widetilde{\alpha}_i\left(U_{1:T}, J, \theta_i; U^*_{1:T}, J^*, \theta^*_i|\theta_{-i}\right) \mathrm{d}U^*_{1:T} d\theta^*_i \\
&\geq \gamma_i^{-1} \int_{A\times B} \tilde{\pi}^N\left(U^*_{1:T}, J^*|\theta_{-i}, \theta^*_i\right) q_{i,1}(\theta^*_i|U_{1:T}, J, \theta_{-i}, \theta_i) \times \\
&\quad \widetilde{\alpha}_i\left(U_{1:T}, J, \theta_i; U^*_{1:T}, J^*, \theta^*_i|\theta_{-i}\right) \mathrm{d}U^*_{1:T} d\theta^*_i \\
&= \gamma_i^{-1} \widetilde{K}_i\left(U_{1:T}, J, \theta_{-i}, \theta_i; A, J^*, \theta_{-i}, B\right),
\end{aligned}$$

which proves (13). Apply (13) repeatedly to get (14). ■

Lemma 5 can be used to find sufficient conditions for the existence of minorization conditions for Sampling Scheme 1 as given in the theorem below, which is similar to Andrieu and Roberts (2009) , Theorem 8. Let $\mathscr{L}_N\{V^{(n)} \in \cdot\}$ denote the sequence of distribution functions of the random variables $\{V^{(n)} : n = 1,2,\dots\}$, generated by Sampling Scheme 1, and let $|\cdot|_{TV}$ be total variation norm of $\cdot$.

Theorem 6 *Suppose that Sampling Scheme 3 satisfies the following minorization condition: there exists a constant $\varepsilon > 0$, a number $n_0 \geq 1$, and a probability measure ν on $\mathscr{V}$ such that $\widetilde{K}^{n_0}(v;A) \geq \varepsilon\,\nu(A)$ for all $v \in \mathscr{V}, A \in \mathscr{B}(\mathscr{V})$. Suppose also that the conditions of Lemma 5 are satisfied. Then Sampling Scheme 1 satisfies the minorization condition*

$$K^{n_0}(v;A) \geq \left(\prod_{i=1}^{p_1} \gamma_i\right)^{-n_0} \varepsilon\,\nu(A)$$

and for all starting values for the Markov Chain

$$\left|\mathscr{L}_N\{V^{(n)} \in \cdot\} - \tilde{\pi}^N\left\{V^{(n)} \in \cdot\right\}\right|_{TV} \leq (1-\delta)^{\lfloor n/n_0 \rfloor},$$

where $0 < \delta < 1$ and $\lfloor n/n_0 \rfloor$ is the greatest integer not exceeding n/n_0.

Proof. To show the first part, suppose $\widetilde{K}^{n_0}(v;A) \geq \varepsilon\,\nu(A)$ for all $v \in \mathscr{V}, A \in \mathscr{B}(\mathscr{V})$. Fix $v \in \mathscr{V}, A \in \mathscr{B}(\mathscr{V})$. Applying Lemma 5 repeatedly gives

$$\begin{aligned}
K^{n_0}(v;A) &\geq \left(\prod_{i=1}^{p_1} \gamma_i\right)^{-n_0} \widetilde{K}^{n_0}(v;A) \\
&\geq \left(\prod_{i=1}^{p_1} \gamma_i\right)^{-n_0} \varepsilon\,\nu(A)
\end{aligned}$$

as required. The second part follows from the first part and Roberts and Rosenthal (2004), Theorem 8. ■

The case when $q_{i,2}(\cdot|U_{1:T}, J, \theta) = \psi(\cdot|\theta)$ is important because it is commonly used for PMMH in the literature, see, for example, Andrieu et al. (2010). Substituting into (12) and using (5) gives

$$\begin{aligned}
\frac{\tilde{\pi}^N\left(U^*_{1:T}|\theta\right)}{q_{i,2}(U^*_{1:T}|U_{1:T}, J, \theta)} &= \frac{\tilde{\pi}^N\left(U^*_{1:T}|\theta\right)}{\psi\left(U^*_{1:T}|\theta\right)} \\
&= \frac{Z(U_{1:T}, \theta)}{p(y_{1:T}|\theta)}. \qquad (15)
\end{aligned}$$

Lemma 7 gives sufficient conditions for Lemma 5 to hold when

$$q_{i,2}(\cdot|U_{1:T}, J, \theta) = \psi(\cdot|\theta).$$

The first condition is from Andrieu et al. (2010).

Lemma 7 *Suppose*

(i) There is a sequence of finite, positive constants $\{c_t : t = 1,\dots,T\}$ *such that for any* $x_{1:t} \in \mathscr{X}_t(\theta)$ *and all* $\theta \in \mathscr{S}$, $f_\theta(x_t|x_{t-1})g_\theta(y_t|x_t) \leq c_t\, m_t^\theta(x_t|x_{t-1})$.
(ii) There exists $\varepsilon > 0$ *such that for all* $\theta \in \mathscr{S}$, $p(y_{1:T}|\theta) > \varepsilon$.

If (i) and (ii) hold then the conditions in Lemma 5 are satisfied.

Proof. Part (i) implies that for all $\theta \in \mathscr{S}$ and all $U_{1:T} \in \mathscr{U}$, $Z(U_{1:T},\theta) \leq \prod_{t=1}^T c_t$. Hence Part (ii) implies that

$$\frac{Z(U_{1:T},\theta)}{p(y_{1:T}|\theta)} < \frac{\prod_{t=1}^T c_t}{\varepsilon}.$$

Substituting into (15) gives the result. ■

Remark 6 *If the state is sampled using backward simulation, similar arguments can be applied to obtain corresponding results (see Appendix B). The mathematical details of the derivation use the results in Olsson and Ryden (2011) and Lindsten and Schön (2012).*

B Backward simulation

Godsill et al. (2004) introduce an alternative way of sampling a particle from a particle filter. Their approach differs from *ancestral tracing* as it samples the indices $J_T, J_{T-1},\dots,J_1$ sequentially, instead of sampling one index J and tracing back its ancestral lineage. The *backward simulation* algorithm has been used in the PMCMC setting by Olsson and Ryden (2011) (in the PMMH algorithm) and Lindsten and Schön (2012) (in the PG algorithm). Chopin and Singh (2013) studied the PG algorithm with *backward simulation* and found that it yields a smaller autocorrelation than the corresponding algorithm using *ancestral tracing*. Moreover, it is more robust to the resampling scheme (multinomial resampling, systematic resampling, residual resampling or stratified resampling) used in the resampling step of the algorithm.

The *backward simulation* algorithm is given below.

Algorithm 3 (Backward Simulation) *Sample* $J_T = j_t$ *conditional on* $u_{1:T}$, *with probability proportional to* $w_T^{j_T}$, *and choose* $x_T^{j_T}$;
2. *For* $t = T-1,\dots,1$, *sample* $J_t = j_t$ *conditional on*

$$(u_{1:t}, j_{t+1:T}, x_{t+1}^{j_{t+1}},\dots,x_T^{j_T}),$$

with probability proportional to $w_t^{j_t} f_\theta(x_{t+1}^{j_{t+1}}|x_t^{j_t})$, *and choose* $x_t^{j_t}$.

We denote the particles selected and the trajectory selected by $x_{1:T}^{j_{1:T}} = (x_1^{j_1},\dots,x_T^{j_T})$ and $j_{1:T}$, respectively. With some abuse of notation, we denote

$$x_{1:T}^{(-j_{1:T})} = \left\{x_1^{(-j_1)},\dots,x_T^{(-j_T)}\right\}.$$

It will simplify the notation to sometimes use the following one-to-one transformation

$$(u_{1:T}, j_{1:T}) \leftrightarrow \left\{x_{1:T}^{j_{1:T}}, j_{1:T}, x_{1:T}^{(-j_{1:T})}, a_{1:T-1}\right\}$$

We will switch between the two transformations and use whichever is more convenient.

The augmented space in this case consists of the particle filter variables $U_{1:T}$ and the sampled trajectory $J_{1:T}$ and PMCMC methods using *backward simulation* target the following density

$$\begin{aligned}\tilde{\pi}_{BSi}^N\left(x_{1:T}, j_{1:T}, x_{1:T}^{(-j_{1:T})}, a_{1:T-1}, \theta\right) :=& \\ \frac{p(x_{1:T},\theta|y_{1:T})}{N^T} \frac{\psi(u_{1:T}|\theta)}{m_1^\theta\left(x_1^{b_1}\right)\prod_{t=2}^T w_{t-1}^{a_{t-1}^i} m_t^\theta\left(x_t^{b_t}|x_{t-1}^{a_{t-1}^{b_t}}\right)} \times& \\ \prod_{t=2}^T \frac{w_t^{a_{t-1}^{j_t}} f(x_t^{j_t}|x_{t-1}^{a_{t-1}^{j_t}})}{\sum_{i=1}^N w_t^{a_{t-1}^i} f(x_t^i|x_{t-1}^{a_{t-1}^i})}.& \end{aligned} \tag{16}$$

Olsson and Ryden (2011) show that, under Assumption 1,

$$\tilde{\pi}^N_{BSi}\left(x_{1:T}, j_{1:T}, x_{1:T}^{(-j_{1:T})}, a_{1:T-1}, \theta\right)$$

has the following marginal distribution

$$\tilde{\pi}^N_{BSi}(x_{1:T}, j_{1:T}, \theta) = \frac{p(x_{1:T}, \theta|y_{1:T})}{N^T}$$

and hence

$$\tilde{\pi}^N_{BSi}(x_{1:T}, \theta) = p(x_{1:T}, \theta|y_{1:T}).$$

Similarly to Lemma 1, Assumption 2 ensures that $\tilde{\pi}^N_{BSi}(u_{1:T}|\theta)$ is absolutely continuous with respect to $\psi(u_{1:T}|\theta)$.

The conditional sequential Monte Carlo algorithm used in the backward simulation also changes. The conditional sequential Monte Carlo algorithm is given in Lindsten et al. (2014) and generates from the full conditional distribution

$$\tilde{\pi}^N_{BSi}\left(x_{1:T}^{(-j_{1:T})}, a_{1:T-1}|x_{1:T}, j_{1:T}, \theta\right).$$

The general sampler using *backward simulation* is analogous to the *ancestral tracing* general sampler, but on an expanded space.

Sampling Scheme 4 (general-BSi) *Given initial values for $U_{1:T}$, $J_{1:T}$ and θ, one iteration of the MCMC involves the following steps*

1. (PMMH sampling) For $i = 1, \ldots, p_1$

Step i:

(a) Sample $\theta_i^* \sim q_{BSi,i,1}(\cdot|U_{1:T}, J_{1:T}, \theta_{-i}, \theta_i)$.

(b) Sample $U_{1:T}^* \sim q_{BSi,i,2}(\cdot|U_{1:T}, J_{1:T}, \theta_{-i}, \theta_i^*)$.

(c) Sample $J_{1:T}^*$ from $\tilde{\pi}^N_{BSi}(\cdot|U_{1:T}^*, \theta_{-i}, \theta_i^*)$.

(d) Accept the proposed values θ_i^*, $U_{1:T}^*$ and $J_{1:T}^*$ with probability

$$\alpha_i(U_{1:T}, J_{1:T}, \theta_i; U_{1:T}^*, J_{1:T}^*, \theta_i^*|\theta_{-i}) = \tag{17}$$
$$1 \wedge \frac{\tilde{\pi}^N_{BSi}(U_{1:T}^*, \theta_i^*|\theta_{-i})}{\tilde{\pi}^N_{BSi}(U_{1:T}, \theta_i|\theta_{-i})} \frac{q_{BSi,i}(U_{1:T}, \theta_i|U_{1:T}^*, J_{1:T}^*, \theta_{-i}, \theta_i^*)}{q_{BSi,i}(U_{1:T}^*, \theta_i^*|U_{1:T}, J_{1:T}, \theta_{-i}, \theta_i)}$$

where

$$q_{Bsi,i}(U_{1:T}^*, \theta_i^*|U_{1:T}, J_{1:T}, \theta_{-i}, \theta_i) =$$
$$q_{BSi,i,1}(\theta_i^*|U_{1:T}, J_{1:T}, \theta_{-i}, \theta_i) q_{BSi,i,2}(U_{1:T}^*|U_{1:T}, J_{1:T}, \theta_{-i}, \theta_i^*).$$

2. (PG or PMwG sampling) For $i = p_1 + 1, \ldots, p$

Step i:

(a) Sample $\theta_i^* \sim q_i(\cdot|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i)$.

(b) Accept the proposed value θ_i^* with probability

$$\alpha_i\left(\theta_i; \theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right)$$
$$= 1 \wedge \frac{\tilde{\pi}^N\left(\theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right)}{\tilde{\pi}^N\left(\theta_i|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}\right)} \frac{q_i(\theta_i|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i^*)}{q_i(\theta_i^*|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i)}.$$

3. Sample $U_{1:T}^{(-J),*} \sim \tilde{\pi}^N(\cdot|X_{1:T}^J, B_{1:T-1}^J, J, \theta_{-i}, \theta_i^*)$.

4. Sample $J \sim \tilde{\pi}^N(\cdot|U_{1:T}, \theta)$.

The PMMH steps in Sampling Scheme 4 simplify in a similar way to Sampling Scheme 1. Olsson and Ryden (2011) show that

$$\frac{\tilde{\pi}^N_{BSi}(U_{1:T}, \theta_i|\theta_{-i})}{\psi(U_{1:T}|\theta_{-i}, \theta_i)} = \frac{Z(U_{1:T}, \theta) p(\theta_i|\theta_{-i})}{p(y_{1:T}|\theta_{-i})},$$

which is the same expression as (5). Hence, if

$$q_{BSi,i,2}(\cdot|U_{1:T}, J, \theta) = \psi(\cdot|\theta)$$

then the Metropolis-Hasting acceptance probability in (17) simplifies to

$$1 \wedge \frac{Z(\theta_i^*, \theta_{-i}, U_{1:T}^*)}{Z(\theta_i, \theta_{-i}, U_{1:T})} \frac{q_{BSi,i,1}(\theta_i|U_{1:T}^*, J^*, \theta_{-i}, \theta_i^*) p(\theta_i^*|\theta_{-i})}{q_{BSi,i,1}(\theta_i^*|U_{1:T}, J, \theta_{-i}, \theta_i) p(\theta_i|\theta_{-i})}.$$

The results in Appendix A can be modified for the distribution $\tilde{\pi}^N_{BSi}(\cdot)$, instead of the distribution $\tilde{\pi}^N(\cdot)$ in a straightforward way. Details are available from the authors on request.

C Smoothness Prior

Suppose that $y_t = g(s_t) + e_t, t = 1,\dots,T$, with the e_t iid $N(0,\sigma^2)$; $g(s)$ is an unknown smooth function, and $0 \le s_1 \le s_2 \le \cdots \le s_T$. We follow Wahba (1978) and place a smoothness prior on $g(s)$ which is defined by the second order differential equation

$$\frac{d^2 g(s)}{ds^2} = \tau \frac{dW(s)}{ds}, \tag{18}$$

where $W(s)$ is a Wiener process with $W(0) = 0$, $\mathrm{Var}(1) = 1$ and τ is a scale parameter. Define $X_t = \left(g(s_t), dg(s_t)/dt\right)'$. We complete the prior specification for g by assuming that X_1 is diffuse, i.e.,

$$p(X_1) \propto 1. \tag{19}$$

It then follows as in Wahba (1978) that $\widehat{g}(s) = E(g(s)|y)$ is a cubic smoothing spline, where $y = (y_1,\dots,y_T)$.

Following Wecker and Ansley (1983), we can write the evolution of X_t for $t = 1,\dots,T-1$ as

$$X_{t+1} = F_t X_t + \eta_t, \quad \eta_t \sim N(0, \tau^2 U_t)$$

where

$$F_t = F(\delta_t) = \begin{pmatrix} 1 & \delta_t \\ 0 & 1 \end{pmatrix}, \quad U_t = \begin{pmatrix} \delta_t^3/3 & \delta_t^2/2 \\ \delta_t^2/2 & \delta \end{pmatrix} \tag{20}$$

with $\delta_t = s_{t+1} - s_t$.